


AD-A242 377



1

**A COMPUTATIONAL MODEL  
OF SEMANTIC MEMORY IMPAIRMENT:  
MODALITY-SPECIFICITY  
AND EMERGENT CATEGORY-SPECIFICITY**  
Technical Report AIP - 147


Martha J. Farah and James L. McClelland

University of Pittsburgh  
Pittsburgh, PA 15260

September, 1991

This research was supported by the Computer Sciences Division, Office of Naval Research, under Contract Number N00014-86-K-0678. Reproduction in whole or in part is permitted for purposes of the United States Government. Approved for public release; distribution unlimited.

In press, *Journal of Experimental Psychology: General*.



91 1106 014

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

## REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION Unclassified		1b RESTRICTIVE MARKINGS	
2a SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release: Distribution unlimited	
2b DECLASSIFICATION / DOWNGRADING SCHEDULE		5. MONITORING ORGANIZATION REPORT NUMBER(S)	
4. PERFORMING ORGANIZATION REPORT NUMBER(S) AIP-147			
6a. NAME OF PERFORMING ORGANIZATION Carnegie-Mellon University	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Computer Sciences Division Office of Naval Research	
6c. ADDRESS (City, State, and ZIP Code) Department of Psychology Pittsburgh, PA 15213		7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, Virginia 22217-5000	
8a. NAME OF FUNDING / SPONSORING ORGANIZATION Same as Monitoring Organization	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-86-K-0678	
8c. ADDRESS (City, State, and ZIP Code)		10. SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO. PROJECT NO. TASK NO. WORK UNIT ACCESSION NO.	
11. TITLE (Include Security Classification) A computational model of semantic memory impairment: Modality-specificity and emergent category-specificity			
12. PERSONAL AUTHOR(S) Farah, Martha J and McClelland, James L.			
13a. TYPE OF REPORT Technical	13b. TIME COVERED FROM TO	14. DATE OF REPORT (Year, Month, Day) 1991, September	15. PAGE COUNT 56
16. SUPPLEMENTARY NOTATION In press, Journal of Experimental Psychology: General			
17. COSATI CODES FIELD GROUP SUB-GROUP		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) Brain damage can cause the selective loss of knowledge about living or nonliving things. This seems to imply that semantic memory is organized taxonomically, with different components specializing in representing knowledge about living and nonliving things. An alternative view of semantic memory is that it is organized by modality, with different components representing information from different sensorimotor channels. In this article we demonstrate how a modality-specific semantic memory system can account for category-specific impairments after brain damage. Specifically, in Experiment 1 we test and confirm the hypothesis (originally put forth by Warrington, McCarthy and Shallice) that visual and functional knowledge play different roles in the representation of living and nonliving things. We then describe a parallel distributed processing model of semantic memory in which knowledge is subdivided by modality into visual and functional components. In Experiment 2 we lesion the model and confirm that damage to visual semantics primarily impairs knowledge of living things, and damage to functional semantics primarily impairs			
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Alan Meyrowitz		22b. TELEPHONE (Include Area Code) (202) 696-4302	22c. OFFICE SYMBOL N00014

A-1

A computational model of semantic memory impairment:  
Modality-specificity and emergent category-specificity

Martha J. Farah and James L. McClelland

Carnegie Mellon University

Address for correspondence:

Department of Psychology  
Carnegie Mellon University  
Pittsburgh, PA 15213 USA

### Abstract

Brain damage can cause the selective loss of knowledge about living or nonliving things. This seems to imply that semantic memory is organized taxonomically, with different components specializing in representing knowledge about living and nonliving things. An alternative view of semantic memory is that it is organized by modality, with different components representing information from different sensorimotor channels. In this article we demonstrate how a modality-specific semantic memory system can account for category-specific impairments after brain damage. Specifically, in Experiment 1 we test and confirm the hypothesis (originally put forth by Warrington, McCarthy and Shallice) that visual and functional knowledge play different roles in the representation of living and nonliving things. We then describe a parallel distributed processing model of semantic memory in which knowledge is subdivided by modality into visual and functional components. In Experiment 2 we lesion the model and confirm that damage to visual semantics primarily impairs knowledge of living things, and damage to functional semantics primarily impairs knowledge of nonliving things. In Experiment 3 we demonstrate that the model accounts naturally for a finding that had appeared problematic for a modality-specific architecture, namely impaired retrieval of functional knowledge about living things. Finally, in Experiment 4 we show how the model can account for a recent observation of impaired knowledge of living things only when knowledge is probed verbally.

How is semantic memory organized? Two general answers to this question have been proposed. One is that semantic memory is organized by taxonomic category, such that different parts of the system represent knowledge about objects from different categories. Alternatively, semantic memory could be subdivided by modality of knowledge, such that one component is responsible for visual information about objects, another for auditory information, and so on.

Patients with selective losses of knowledge following brain damage appear to provide a direct source of evidence on the organization of semantic memory. Unfortunately, this evidence yields conflicting answers. In most cases, the losses appear to be tied to specific modalities, resulting in impaired recognition of objects in just one modality (e.g. visual or auditory agnosia) or impaired manipulation of objects with specific uses, despite intact recognition of them (apraxia, in which, e.g., a key might be pulled rather than turned). These observations are consistent with recent neurophysiological data showing that most cortical neurons are modality-specific, even in regions that were traditionally viewed as supramodal association areas (e.g., Sereno & Allman, 1990). In some cases, however, brain damage seems to cause category-specific losses of knowledge, which cut across different modalities. Specifically, there are patients who seem to have lost their knowledge of living things, and others who seem to have lost their knowledge of nonliving things. These observations suggest that the architecture of semantic memory incorporates at least two general, taxonomically-defined subsystems, for representing knowledge of living and nonliving things.

In this article we attempt to resolve the apparent conflict between these two types of neuropsychological evidence. After reviewing the neuropsychological evidence for category-specificity in semantic memory, we will present a parallel distributed processing model whose architecture distinguishes only between modalities of knowledge, but which, when damaged, displays category-specificity similar to that of the patients described in the neuropsychological literature.

#### Impairments in knowledge of living and nonliving things

The most commonly observed semantic memory dissociation is between

impaired knowledge of living things with relatively preserved knowledge of nonliving things. In the first report of this phenomenon, Warrington and Shallice (1984) described four patients who were much worse at identifying living things (animals, plants) than nonliving things (inanimate objects). All four of these patients had recovered from Herpes encephalitis, and all had sustained bilateral temporal lobe damage. Two of the patients were studied in detail, and showed a selective impairment for living things across a range of tasks, both visual and verbal. Table 1 shows examples of their performance in a visual identification task, in which they were to identify by name or description the item shown in a colored picture, and in a verbal definition task, in which the names of these same items were presented auditorily, and they were to define them. Examples of their definitions are also shown in Table 1.

-----  
 Insert Table 1 about here  
 -----

Farah, McMullen and Meyer (1991) studied two head-injured patients whose knowledge of living things appeared to be selectively disrupted. We examined their picture recognition performance as a function of the living/nonliving distinction, as well as many other, possibly confounded, factors that might influence performance, including complexity, familiarity, name frequency, name specificity (i.e., basic object level or subordinate level), and similarity to other objects. A regression analysis showed that, even with all of these factors accounted for, the living/nonliving distinction was an important predictor of recognition performance.

Other cases of selective impairment in knowledge of living things include additional postencephalitic patients described by Pietrini, Nertempi, Revello, Pinna and Ferro-Milone (1988), Sartori and Job (1988), and Silveri and Gianotti (1988), a patient with encephalitis and strokes described by Mehta and Newcombe (1990), and a patient with a focal degenerative disease described by Basso, Capitani and Laiaccona (1988). In all of these cases there was damage to the temporal regions, known to be bilateral except in Farah et al.'s case 2, Pietrini et al.'s

case 1, and the case of Basso et al., where there was evidence only of left temporal damage.

The opposite dissociation, namely impaired knowledge of nonliving things with relatively preserved knowledge of living things, has also been observed. Warrington and McCarthy (1983, 1987) described two cases of global dysphasia following large left hemisphere strokes in which semantic knowledge was tested in a series of matching tasks. Table 2 shows the results of a matching task in which the subjects were asked to point to the picture, in an array, that corresponded to a spoken word. Their performance with animals and flowers was reliably better than with nonliving things. One of these subjects was also tested with a completely nonverbal matching task, in which different-looking depictions of objects or animals were to be matched to one another in an array, and showed the same selective preservation of knowledge of animals relative to inanimate objects.

-----  
 Insert Table 2 about here  
 -----

#### Implications of the living-nonliving dissociations for models of normal semantic memory

The most straightforward interpretation of the double dissociation between knowledge of living and nonliving things is that these two bodies of knowledge are represented by two separate category-specific components of semantic memory. This interpretation is consistent with the view that semantic memory is organized along taxonomic lines, at least as far as the distinction between living and nonliving things is concerned. However, Warrington and colleagues have suggested an alternative interpretation, according to which semantic memory is fundamentally modality-specific. They argue that selective deficits in knowledge of living and nonliving things may reflect the differential weighting of information from different sensorimotor channels in representing knowledge about these two categories. Warrington and McCarthy (1983) and Warrington and Shallice (1984) have pointed out that living things are distinguished primarily by their sensory attributes, whereas nonliving things are distinguished primarily by



their functional attributes. For example, our knowledge of an animal such as a leopard, by which we distinguish it from other similar creatures, is predominantly visual. In contrast, our knowledge of a desk, by which we distinguish it from other furniture, is predominantly functional (i.e., what it is used for.) Thus, the distinctions between impaired and preserved knowledge in the cases reviewed earlier may not be "living/nonliving" distinctions per se, but "sensory/functional" distinctions.

The sensory/functional hypothesis seems preferable to a strict living/nonliving hypothesis for two reasons. First, it is more consistent with what is already known about brain organization. As mentioned earlier, it is well known that different brain areas are dedicated to representing information from specific sensory and motor channels. Functional knowledge could conceivably be tied to the motor system. In any case, there is prior evidence for the selective vulnerability of knowledge of functional attributes following left hemisphere damage: Goodglass and Baker (1976) found that left hemisphere-damaged aphasic patients had particular difficulty relating a named object to a word describing its use, compared to words describing its sensory qualities or words denoting other objects in the same category. A second reason for preferring the sensory/functional hypothesis to the living/nonliving hypothesis is that exceptions to the living/nonliving distinction have been observed in certain cases. For example, Warrington and Shallice (1984) report that their patients, who were deficient in their knowledge of living things, also had impaired knowledge of gemstones and fabrics. Warrington and McCarthy's (1987) patient, whose knowledge of most nonliving things was impaired, seemed to have retained good knowledge of very large outdoor objects such as bridges or windmills. It is at least possible that our knowledge of these aberrant categories of nonliving things is primarily visual.

Unfortunately, there is a problem with the hypothesis that "living things impairments" are just impairments in sensory knowledge, and "nonliving things impairments" are just impairments in functional knowledge. This hypothesis seems to predict that cases of "living things impairment" should show good knowledge of the functional attributes of living things, and cases of "nonliving things impairment"

should show good knowledge of the visual attributes of nonliving things. The evidence available in cases of "nonliving things impairment" is limited to performance in matching-to-sample tasks, which does not allow us to distinguish knowledge of visual or sensory attributes from knowledge of functional attributes. However, there does appear to be adequate evidence available in cases of "living things impairment," and in at least some cases it disconfirms these predictions.

Knowledge of nonvisual attributes of living things in cases of "living things impairment"

Consider the definitions of living and nonliving things given by Warrington and Shallice's (1984) two cases (Table 1). Although the definitions of nonliving things may be somewhat skimpy on visual detail, in keeping with the sensory/functional hypothesis, the definitions of living things do not show preserved functional knowledge. If these cases have lost just their visual semantic memory, they should be able to retrieve the functional attributes of living things, for example the fact that parrots are kept as pets and can talk, that daffodils are a spring flower, and so on.

In the other cases of "living things impairment," visual and functional knowledge have been compared directly, and functional knowledge of living things ranges from mildly to severely impaired. Mehta and Newcombe (1990) presented their subject with triads of words, with the instruction to group together two of the words according to either the visual similarity of the words' referents or some factual commonality (e.g., normally found in the UK). When the words named nonliving things, their case performed within normal limits. However, when the words named living things, their case performed significantly worse than control subjects, even when the grouping was based on factual rather than visual properties. Silveri and Gianotti (1988) assessed the ability of their patient to identify animals on the basis of two kinds of spoken definition: Visual descriptions of the animal's appearance, such as "an insect with broad, colored, ornate wings" for "butterfly", and nonvisual descriptions, of either metaphorical verbal associations to the animal, such as "king of the jungle" for "lion," or functions of the animal, such as "the farm animal which bellows and supplies us with

milk" for "cow." Although the subject was worse at identifying animals from visual descriptions than from nonvisual descriptions, he performed poorly with both, and identified only 58% of the animals on the basis of nonvisual descriptions (which control subjects had rated "easy"). Basso et al.'s (1988) patient also appeared to be better at retrieving nonvisual information about living things than he was at retrieving visual information, but nonvisual information was not intact. These different types of knowledge were tested by naming a word, and then asking a multiple-choice question about it, tapping categorical information such as "is it a bird, mammal, fish or reptile?", functional information such as "does it live in Italy or the desert?," or visual information such as "does it have a smooth back or is it hump-backed?" The patient performed at chance on the categorical as well as the visual questions, and performed less than perfectly with the functional questions (35/42). (It should be noted that not all of the words denoted living things. The authors tested the patient with words he had failed to match with pictures; most of these were living things. The authors did not separately report the results for living and nonliving things.)

Similarly, Sartori and Job (1988) found better performance in tests tapping nonvisual than visual knowledge of animals in their case, but their subject nevertheless appeared mildly impaired in nonvisual tasks. For example, in defining living and nonliving things, the subject made numerous factual errors about nonvisual characteristics of animals and vegetables, twice as many as he made about nonliving things. He also made occasional errors in identifying animals with their characteristic sounds or environments, although in the absence of normative data it is difficult to interpret these results. Farah, Hammond, Mehta and Ratcliff (1989) tested the ability of one of the head-injured patients described earlier (Farah et al., 1991, case 1) to retrieve visual and nonvisual knowledge about living and nonliving things, and compared his performance to age- and education-matched normal subjects. We found that his performance fell outside of normal limits only for visual knowledge of living things. However, whereas he performed at an average level in retrieving nonvisual information about nonliving things, he performed below average in retrieving nonvisual

information about living things, and the discrepancy between his performance with these two kinds of question was larger than for any of the twelve control subjects. Our unpublished observations of case 2 are that she was impaired at retrieving functional information about animals, such as knowing which animal provides wool, as well as at recognizing animal sounds. When given the test designed by Farah et al. (1989), she performed at chance on the questions concerning visual as well as nonvisual properties of living things, whereas she performed far above chance with nonvisual properties of nonliving things.

In sum, the sensory/functional hypothesis seems more attractive than the living/nonliving hypothesis because it is more in keeping with what we already know about brain organization. However, it does not seem able to account for all of the data. In particular, it does not seem able to account for the impaired ability of these patients to retrieve nonvisual information about living things.

The goal of our model is to demonstrate that the sensory/functional hypothesis is sufficient to account for these semantic memory impairments when it is taken together with a certain conception of mental representation; specifically, the idea of active, distributed representations, in which the activation of the representation depends on mutual support among different parts of the representation. This idea is common to a wide range of recurrent parallel distributed processing (PDP) models (e.g., Anderson, Silverstein, Ritz & Jones, 1977; McClelland & Rumelhart, 1985). We will show that a model of semantic memory with active distributed representations consisting of just two types of semantic information, visual and functional, can be lesioned to produce selective impairments in knowledge of living things and nonliving things. More importantly, we will show how such a model can account naturally for the impairment of both visual and functional knowledge of living things following damage confined to visual semantics. Finally, we will also show how this model can account for a recently described case in which knowledge of living things was impaired only when probed verbally, which had initially been interpreted as evidence that semantic memory is subdivided not only by category of knowledge but also by modality of access.

Before presenting the simulation model, and the results of lesioning the model, we will describe an experiment that tests the basic assumption of the sensory/functional hypothesis, namely that living things are known primarily by their sensory features, and that nonliving things are known primarily by their functional features.

### Experiment 1

In this experiment, normal subjects read dictionary definitions of living and nonliving things, and underlined all occurrences of visual and functional descriptors. This tested whether there is a difference in the importance of sensory (specifically, visual) and functional properties for the meaning of living and nonliving things, and provided us with a quantitative estimate of the ratio of visual to functional features for the representations of living and nonliving things in the model.

### Methods

Materials. The lists of living and nonliving things were taken from Warrington and Shallice's (1984) Experiment 2. Definitions were copied from the American Heritage Dictionary, and printed in a random order.

Procedure. Subjects either read for visual descriptors or functional descriptors. If they read for visual descriptors, they were told to underline all occurrences of words describing any aspect of the visual appearance of an item. If they read for functional descriptors, they were told to underline all occurrences of words describing what the item does, or what it is for.

Subjects. Forty-two undergraduate students from Carnegie Mellon University participated in exchange for course credit. Half read for visual descriptors, and half for functional descriptors.

### Results and Discussion

Subjects who read for visual descriptors underlined an average of 2.68 visual descriptors for each living thing, and 1.57 for each nonliving thing. Subjects who read for functional descriptors underlined an average of 0.35 functional descriptors for each living

thing, and 1.11 for each nonliving thing. The resultant ratios of visual to functional features for living things is 7.7:1, and for nonliving things is 1.4:1. Thus, these data confirm the hypothesis that visual attributes are more important than functional attributes for defining living things, but do not support the converse hypothesis that functional attributes are more important than visual attributes for defining nonliving things: Subjects found more visual descriptors than functional descriptors in the definitions of nonliving things, but did not find more functional than visual descriptors for nonliving things. One of the interesting conclusions of the simulation to be described is that a large difference in the number of visual and functional attributes for living things, with a much smaller difference in the same direction for nonliving things, is sufficient to account for both the "living things" impairments and the "nonliving things" impairments. The overall ratio of visual to functional features, combining living and nonliving things, is 2.9:1.

#### Model

In parallel distributed processing (PDP) systems, a representation consists of a pattern of activation across a network of highly interconnected neuron-like units (Anderson, Silverstein, Ritz & Jones, 1977; Hinton, McClelland & Rumelhart, 1986; McClelland & Rumelhart, 1985). The units can be thought of as each representing some aspect of the entity being represented by the pattern (although these aspects need not be nameable features, or correspond in any simple way to our intuitions about the featural decomposition of these concepts). For example, in the case of living and nonliving things, some of the units would represent aspects of the visual qualities of the item, and other units would represent aspects of the item's functional roles. The extent to which activation in one unit causes activation in the other units to which it is connected depends upon the connection strengths, or "weights," between the units. Presenting a stimulus to the network results in an initial pattern of activation across the units, with some units being activated and others not. This pattern will then begin to change as each unit receives activation from the other units to which it is connected within the network. Eventually a stable pattern will

result, with each unit holding a particular activation value as a result of the inputs it is receiving from the other units to which it is connected.<sup>1</sup>

Figure 1 shows the architecture of the model. There are three main pools of units, corresponding to verbal inputs or outputs (name units), visual inputs or outputs (picture units), and semantic memory representations. The semantic memory units are divided into visual units and functional units. There are bidirectional connections between units both within and between pools, with the exception that there are no direct connections between the name and picture units. There are 24 name units, 24 picture units, and 80 semantic memory units, divided into 60 visual semantic and 20 functional semantic units according to the roughly 3:1 ratio obtained in Experiment 1.

-----  
Insert Figure 1 about here  
-----

The specific processing assumptions of this model are the same as for the distributed memory model of McClelland and Rumelhart (1985). In brief, units can take on continuous activation values between -1 and +1. The weights on the connections between units can take on any real values (positive, negative, or zero). There are no thresholds in the model, and the influence of each unit on the input to each other unit is just the activation of the influencing unit times the strength of the relevant connection. Processing is synchronous, that is, on each cycle the total input to each unit is calculated on the basis of the activation levels of the units to which it is connected and the weights on those connections, and the activation levels of all units are then updated simultaneously. Activation levels are updated according to a nonlinear activation function, which keeps activations bounded between -1 and 1. Inputs are presented for 10 cycles.

---

<sup>1</sup>Many recent connectionist models, such as the models of spelling-to-sound translation of Rosenberg & Sejnowski (198 ) or Seidenberg & McClelland (1989), have employed only unidirectional connections from input via internal units to output, and have computed activations in a single, feed-forward pass. At least in the latter case, the use of a feed-forward architecture was a simplification adopted for the sake of tractability, and did not represent a change of principle in favor feedforward information processing.

Ten living and 10 nonliving things were represented as randomly generated patterns of -1's and +1's over all three pools of units. The representation of each item included the full 24 name units and picture units, but only subsets of the semantic memory units in order to capture the different ratios of visual and functional information in living and nonliving things. Living things were represented with an average of 16.1 visual and 2.1 functional units, and nonliving things were represented with an average of 9.4 visual and 6.7 functional units. All patterns contained both types of semantic memory unit.

A simple error-correcting learning procedure was used to train the network to produce the correct semantic and name pattern when presented with each picture pattern, and the correct semantic and picture pattern when presented with each name pattern. On each training trial, the name or the picture corresponding to one of the living or nonliving things was presented to the name or picture input units, and the network was allowed to settle for 10 cycles. The weights among the units were then adjusted using the delta rule (Rumelhart, Hinton & McClelland, 1986) to minimize the difference between the resultant activation of each unit and its correct activation.<sup>2</sup> To distribute the work of producing the desired outputs over as much of the network as possible, the weights were all multiplicatively reduced by 2% of their value at the end of each training epoch (i.e., each pass through the full set of 40 training trials). This procedure, known as "weight decay," tends to keep individual weights from growing large, thereby forcing the network to distribute the associations across a larger number of connections. This results in networks that are more resistant to partial damage. Training was continued for 100 epochs. From the point of view of the training procedure there are no hidden units, so back-propagation is not necessary.

In order to assess the generality of results obtained with this model, four variants of the model were also tested. The first two variants consisted of the exact same architecture and training procedure

---

<sup>2</sup>Note that the training procedure is not meant to simulate the process by which people acquire semantic memory knowledge. It is merely a tool for creating a pattern of connection strengths that embodies the assumed associations between patterns in the different pools of units.



for setting weights, but with training terminated after 50 and 200 epochs. Although both of these variants were trained sufficiently well that they performed perfectly before damage, the different final patterns of weights might be expected to respond differently to damage. The third variant consisted of the same architecture with a different training procedure. In this case, there was no weight decay, and the network was therefore expected to show less resistance to damage. A fourth variant consisted of the original architecture and training procedure, but with a different proportion of visual and functional semantic units in the model. Because one group of subjects in Experiment 1 identified visual attributes used in defining living and nonliving things, and the other group identified functional attributes used in defining living and nonliving things, the ratio of visual to functional semantic units obtained in Experiment 1 was computed from different subjects' data. Instead of using the results of Experiment 1 to set this ratio in the model, in the third variant we arbitrarily set the numbers of visual and functional semantic units to be equal, in other words 40 semantic units of each type. We used the data of Experiment 1 only to set the ratios of visual units in the representations of living and nonliving things, and of functional units in the representations of living and nonliving things, which were ratios obtained within subjects. In this version of the model, living things were represented with an average of 10.6 visual and 4.0 functional units, and nonliving things were represented with an average of 6.2 visual and 12.8 functional units. The effects of lesions on the performance of the basic model and its variants were then explored.

#### Experiment 2.

The goal of this experiment is to test the hypothesis that selective impairments in knowledge of living and nonliving things can be explained by selective damage to visual and functional semantic memory representations, respectively. We test this hypothesis by lesioning the model and observing its performance at associating pictures and names of both living and nonliving things. Picture-naming is a kind of picture-name association task, in which the picture is given and the name must be produced. In this model, picture naming consists of presenting the

picture portion of a pattern in the picture units, letting the network settle, and then reading the resultant pattern in the name units. Matching to sample, as used by Warrington and McCarthy (1983, 1987) is another kind of picture-name association task, in which the name is given, and the correct picture must be selected from amongst a choice set. In this model, it consists of presenting the name portion of a pattern in the name units, letting the network settle, and then reading the resultant pattern in the picture units. In each case, the model's performance on each pattern was scored as correct if the resulting pattern matched the correct pattern more closely than any of the other 19 possible patterns.

### Methods

Procedure. Twelve types of simulation were run, corresponding to 0%, 20%, 40%, 60%, 80% and 99% damage to the visual and to the functional semantic memory units. The different degrees of damage were brought about by subjecting each unit of the relevant pool of semantic memory units to a 0, .2, .4, .6, .8. or .99 chance of being damaged. Each of the twelve simulations was damaged 5 times each, with the damage being re-applied to an intact network each time. For each of these simulations, forty picture-name association trials were run: twenty picture-naming trials, in which each of the picture patterns was presented to the network and the resultant name patterns were scored, and twenty matching-to-sample trials, in which each of the name patterns was presented to the network and the resultant picture patterns scored. This procedure was applied to the original model, and to the four variants described earlier.

### Results and Discussion

Table 3 shows the results from the simulations of visual and functional semantic memory damage to the basic model. When visual semantic memory units are damaged, the effect is greater on the naming of living things than nonliving things. As can be seen in Figure 2, the greater the damage, the greater the dissociation between performance with living and nonliving things. When functional semantic memory units

are damaged, the only effect is on nonliving things, and this effect also increases with increasing damage.

-----  
 Insert Table 3 about here  
 -----

-----  
 Insert Figure 2 about here  
 -----

-----  
 Insert Figures 3a-d about here  
 -----

The pattern of results obtained with the four variants of the model were similar, as shown in Figure 3a-d. Figure 3a show the results of visual and functional semantic memory damage when learning was terminated after half as many trials as in the basic model. Figure 3b show the results of lesioning visual and functional semantic memory when learning continued for twice as long as in the basic model. In both variants, visual semantic memory damage affects performance with living things more than with nonliving things, and functional semantic memory damage affects performance with nonliving things more than with living things. Figure 3c show the effects of semantic memory damage on the model trained without weight decay. Damage has a much larger effect overall on the performance of this model, consistent with the tendency of weight decay to produce more distributed and thus more robust representations. However, as in the previous models, damage to visual semantics impairs performance on living things more than on nonliving things, and damage to functional semantics has the opposite effect. Figure 3d show the results of lesioning a model in which the overall numbers of visual and functional semantic memory units were arbitrarily set to be equal, with the ratios of each type of semantic memory attribute in the representations of items being set by the within-subject data from Experiment 1, as before. As in the previous models, lesioning visual semantics causes disproportionate impairment of performance with living things, and lesioning functional semantics causes disproportionate impairment of performance with nonliving things.

Another way of assessing the effects of damage to either visual or functional semantics on the network's knowledge of living and nonliving things is to compare the pattern of activation obtained in the semantic units after damage when a picture or name is presented to that obtained before damage. One way to quantify this comparison is using the dot product of the pattern obtained and the target pattern. The bigger the dot product, the better the match. Table 4 and Figure 4 show the average dot products, normalized to 1 for the undamaged network, for the semantic memory patterns after different degrees of damage to visual and functional semantics for the basic model. Figures 5a-d shows the same information graphically for the four variants of the basic model. The dot products indicate that damage to visual semantics impairs the semantic representation of living things more than nonliving things, and damage to functional semantics has the opposite effect.

-----  
 Insert Table 4 about here  
 -----

-----  
 Insert Figure 4 about here  
 -----

-----  
 Insert Figures 5a-d about here  
 -----

In summary, the basic prediction of the sensory/functional hypothesis was borne out: Damage to visual semantic memory impaired knowledge of living things to a greater extent than nonliving things, and damage to functional semantic memory impaired knowledge of nonliving things to a greater extent than living things. This result was general across five different implementations of the model, and across two different ways of measuring model performance.

### Experiment 3.

Earlier it was noted that at least some cases of "living things impairment" are impaired at accessing functional, as well as visual, information about living things. On the face of things, this phenomenon

seems to disconfirm the sensory/functional hypothesis, and require that the model incorporate into its architecture an explicit distinction between knowledge of living and nonliving things. The goal of this experiment is to find out whether the model can account for impaired access to functional information about living things following damage to visual semantic memory units.

If it were the case that representations need a certain "critical mass" to become activated, so that even if a portion of the representation were spared by brain damage it could not be accessed in the absence of other parts of the representation, then the sensory/functional hypothesis could explain the apparent across-the-board impairments in knowledge of living things as follows: Given that most of the semantic memory features in the representations of living things are visual features, and they have been destroyed, then those few functional features associated with the representation might lack the critical mass to become activated. In fact, most parallel distributed processing models display just this critical mass effect. It arises because the ability of any given unit to attain and hold its proper activation value depends upon collateral connections with other units in the network. Although PDP systems are robust to small amounts of damage, if a large proportion of the units participating in a given representation are destroyed, the remaining units will not receive the necessary collateral inputs to achieve their proper activation values.

Procedure. Rather than elaborate the model with additional pools of input and output units to represent questions and answers, for the purpose of simulating question-answering tasks, we have assessed the availability of functional semantic memory information in the model directly: Input patterns (names or pictures) were presented, the network was allowed to settle, and the resultant patterns of activation in the functional semantic memory units were recorded. As in the previous experiment, the quality of the semantic memory representation was measured by a normalized dot product, in this case in just the functional semantic memory units. The procedures for training and damaging the model were the same as for Experiment 2.

### Results and Discussion

Table 5 and Figure 6 show the average scaled dot products of the obtained and correct functional semantic memory patterns for living and nonliving things at each degree of damage to the visual semantic memory units. As predicted, damage to visual semantic memory impairs access to functional semantic memory disproportionately for living things. As can be seen in Figures 7a-d, essentially the same results were obtained for the four variants of the basic model described earlier. The different variants display the effect to different degrees, but all show the same qualitative pattern, namely, impaired activation of functional semantic memory, more so for living than nonliving things, after visual semantic memory damage.

-----  
Insert Table 5 about here  
-----

-----  
Insert Figure 6 about here  
-----

-----  
Insert Figures 7a-d about here  
-----

Although damage to visual semantic memory impairs retrieval of functional knowledge of living things, it affects functional knowledge of living things less than visual knowledge. This can be seen by comparing Figures 4 and 5, which show the dot products of the obtained and correct pattern over all of semantics after visual semantic damage, to Figures 6 and 7, which show the dot products for functional semantics in particular. This pattern is consistent with the behavior of the patients reviewed earlier, whose impairments in knowledge of living things tend to be more obvious in the visual than in the functional domain.

### Experiment 4

A third type of dissociation involving living and nonliving things was recently described by McCarthy and Warrington (1988). They describe a patient with progressive aphasia and left temporal hypometabolism of

unstated etiology. This subject's knowledge of living things appeared to be impaired only when tested verbally. As shown in Table 6, he was able to identify pictures of both living and nonliving things, and he was able to define nonliving things that were named aloud to him. However, he was impaired at defining living things that were named aloud. Table 6 also shows examples of his responses to visually and verbally probed animals.

-----  
 Insert Table 6 about here  
 -----

In their discussion of this patient, McCarthy and Warrington suggest that the pattern of impaired and preserved performance implies that semantic memory may be subdivided by both category and modality of access. According to this interpretation, there is one store of knowledge about living things for access by verbal systems, another store of knowledge about living things for access by visual systems, a store of knowledge about nonliving things for verbal access, and so on. The goal of this experiment was to simulate the behavior of McCarthy and Warrington's (1988) case with the present model, which has neither separate knowledge stores for living and nonliving things, nor for different input modalities. This was accomplished by damaging the connections between the name units and the visual semantics units.

Procedure. The model was damaged by destroying the connections that go from the name units to the visual semantic memory units. Six different simulations were run, corresponding to different degrees of damage to these connections: destruction of 0%, 20%, 40%, 60%, 80% and 100% of the connections between name and visual semantic units, randomly chosen. As in Experiment 2, the performance of the network after damage was tested in two ways. First, we scored the percentage of trials on which, given a picture the correct name could be selected, or given a name the correct picture could be selected. Second, we calculated the normalized dot product between the obtained and target semantic memory patterns when either a picture or a name was presented.

#### Results and Discussion

Table 7 and Figure 8 show the percent correct for name-picture association after different degrees of damage to connections from name units to visual semantics units in the basic model. Like the case of McCarthy and Warrington (1988), the impairment of the model has both category specificity and modality specificity. The model is by far the most impaired with living things presented verbally, next most impaired with nonliving things presented verbally, and least impaired with pictures of either living or nonliving things. One curious aspect of the model's performance is the better comprehension of the names of nonliving things when the connections between names and visual semantics are entirely destroyed than when they are 80% destroyed. The poor performance at 80% disconnection is interpretable as a kind of interference caused by the extremely noisy patterns of activation entering the semantics units from the name units. The 20% remaining connections evidently produce inappropriate patterns of activation in the visual semantics units, thereby interfering with the ability of collateral connections from functional semantics to activate the correct patterns in visual semantics.

-----  
 Insert Table 7 about here  
 -----

-----  
 Insert Figure 8 about here  
 -----

-----  
 Insert Figures 9a-d about here  
 -----

Figures 9a-d show the performance of the four variants of the model when damaged and then tested as described above. The same qualitative pattern of results is found in each case, with the poorest performance by far found for named living things.

Table 8 and Figure 10 show the average normalized dot products of the obtained and correct semantic memory patterns for living and nonliving things presented as names and pictures, for the basic model. Figures 11a-d show the same measures for the four other versions of the



model. The dot products reveal essentially the same qualitative pattern of performance as the percent correct measure. The activation of semantic memory by pictures is relatively unimpaired at all levels of damage in this model, whereas the activation of semantic memory by names is impaired, particularly for the names of living things.

In summary, we have shown that the behavior of McCarthy and Warrington's patient can be accounted for in a relatively parsimonious way, by postulating damaged connections between name units and visual semantics units. One possible objection to this account is based on McCarthy and Warrington's observation that the patient's performance was consistent, in terms of specific items failed, from testing session to testing session. It has been proposed (Shallice, 1987) that consistency implies damage to representations, whereas impaired access to representations should lead to variable performance. It is certainly true that some types of access disorders would lead to variable performance (e.g., noise in a telephone line). However, in the context of the present model, it can be seen that there is no necessary relation between disorders of representation versus access, on the one hand, and damage to units versus connections, on the other. Damage to connections in this model leads to high consistency in items failed. This is because certain connections are more important for activating some representations than others, and so whenever a given subset of connections is destroyed, the subset of representations that is most dependent upon those connections will always suffer.

#### General Discussion

The existence of selective impairments for knowledge of living and nonliving things would seem to imply that the architecture of semantic memory consists of at least some taxonomically-defined components. However, we have shown that a simple model of semantic memory with only modality-specific components can account for all three types of category-specific semantic memory impairment that have been observed with patients. Let us examine some of the general implications of these findings for cognitive psychology and neuropsychology, as well as some

cautions that should be borne in mind while interpreting the results of our model.

#### Limitations of the present model

The model we have presented here is a simple one, designed to test some very general principles concerning the relations between modality-specific and category-specific knowledge. Our goal was to determine whether these principles could account for certain general findings that have emerged across a number of different studies of patients with different impairments in semantic memory. We have not attempted to provide a detailed account of the ways that semantic memory is used in naming pictures, defining words, and so on, nor of the precise nature of the damage in cases of semantic memory impairment.

For example, the model has only two kinds of semantic memory representations: visual and functional. We could have added semantics derived from other perceptual modalities (e.g., auditory, tactile), and we could have subdivided the fairly general concept of "functional" semantic memory into more specific components. Whereas such elaborations of the model might change the sizes of the dissociations found here, they would probably not change the basic qualitative patterns (unless the proportions of added semantic units were negatively correlated with the visual and functional units in terms of the numbers participating in the representations of living and nonliving things).

Another way in which the model is simplified and unrealistic is that there is no difference between name and picture representations in the kinds of relations they have with semantic memory. For example, we might expect that the perceptual representations of pictures would have a closer (more systematic and/or more robust) set of connections with the visual semantic representations than the name representations have. If we had included this difference in the model, we might have found differences between the size of the dissociation found in picture naming compared to purely verbal tasks such as definitions. Specifically, one might expect the effects of damage to visual semantics to be more pronounced in tasks involving picture processing. There is a hint of such a difference in the data from patients shown in Table 1.

For simplicity's sake, we have also assumed that the effects of brain damage can be simulated by destroying the neuron-like units or the connections between such units. However, the effects of Herpes encephalitis, head injury, and stroke on neural functioning may be more fully captured by the combined effects of destroying units and connections, as well as by other changes to the network such as adding noise to the connection strengths or to the activation levels of the units, changing the maximal activation values of the units, or changing the rate at which activation decays. These different ways of damaging the network would be expected to have slightly different effects on its performance after damage. For example, adding noise to a certain pool of units would lead to low consistency in the particular test items failed from one test to another, whereas destroying units or connections would lead to high consistency. Nevertheless, these differences would not change the basic patterns concerning the category-specificity and modality-specificity of the deficits reported here.

A final word of caution in relating our simple model to patient behavior is that the measures of performance that we have used with the model are not the same as those that have been used with patients. The 20-alternative forced choice picture-name association task is somewhat similar to the picture-naming and matching-to-sample tasks that have been used with patients, but reading the dot product of the actual and expected semantic memory patterns is quite an abstraction from the question-answering tasks used with patients! This problem is not unique to comparisons between computer simulations and patients, however. Different patients have been studied with different tasks, which makes precise inter-patient comparisons impossible as well. However, neither the difficulties with precise inter-patient nor simulation-patient comparisons prevents us from generalizing about common qualitative patterns of impairment, and their possible underlying causes.

We also wish to note that the present model is not intended to account for category-specific impairments in cognitive systems other than semantic memory. Selective dissociations have been documented within the visual recognition system, affecting just face recognition or just printed word recognition (e.g., Farah, 1991), and within the lexical system, affecting name retrieval for categories as specific as

colors, letters or body parts (e.g., Goodglass, Wingfield, Hyde & Theurkauf, 1986). From the point of view of the present model, these impairments would be located in the "visual" and "verbal" input systems, which we have not attempted to model with any verisimilitude. Our results are relevant to these other category-specific phenomena only in a very general way: They alert us to the fact that every neuropsychological dissociation need not have a corresponding distinction in the cognitive architecture.

#### General implications

Having enumerated some of the ways in which the present model may be incomplete or inaccurate in detail, and some neuropsychological phenomena which it is not intended to explain, let us review the general principles that the model has been successful in demonstrating. First, the model has shown how category-specific impairments can arise after damage to a system that has no category-specific components. Specifically, it has shown how impairments in knowledge of living things and nonliving things, and even impairments in knowledge of living things when just probed verbally, can be accounted for without postulating a semantic memory system with any inherently category-specific components. Instead, these impairments can all be accounted for by a relatively simple semantic memory architecture, in which there are just two components of semantic memory, which differ from one another by modality and not by category.

The ability of a modality-specific semantic memory architecture to account for category-specific semantic memory impairments depends, of course, on there being a correlation between modality of knowledge and category of knowledge. In this case, it depends on the fact that living things are known by us primarily through their visual attributes, which was suggested years ago by Warrington and her colleagues, and which we verified in Experiment 1. One way of describing the relation between the living/nonliving distinction and the visual/functional distinction is that they are "confounded," in the same way that we might speak of confounded factors in an experiment. However, such a description does not fully capture the degree to which the impairments are category-specific. In patients with impaired knowledge of living things,

knowledge about functional properties of living things is also impaired. This is true of the model as well, and can be explained in terms of a very general property of distributed representations, in which the different parts of the representation provide mutual support for one another. Although such representations are robust to small amounts of damage, larger amounts will deprive the intact parts of the representation of needed support. As a result, even those intact parts will be unable to attain their proper activation levels. Thus, category-specificity is an emergent property of the network under certain kinds of damage.

Figure 1 is a "box-and-arrow" outline of our model, showing the different types of representations involved in semantic memory, and their relations to one another. This is the level of description at which most models are cast in cognitive neuropsychology. In many cases, this level of detail has been sufficient, and many cognitive impairments have been successfully interpreted as the simple deletion of a box or an arrow. However, the semantic memory impairments discussed here provide an example of the limitations of this approach, and of the need to understand what goes on within the boxes. As discussed earlier, it is not apparent why damage to the visual semantic memory component of the model would result in impaired access to functional semantic memory knowledge about living things. To explain this, in the context of the model shown in Figure 1 at any rate, one must describe the system at a more detailed level of analysis, which includes the internal workings of the boxes. The effect of visual semantic memory damage on functional knowledge of living things can be explained in terms of the kinds of representations and computations taking place inside the outlined components in Figure 1. In more general terms, the macrostructure of the system's behavior -- what categories or modalities of knowledge are spared or impaired -- does not just depend upon the macrostructure of the system -- for example, what different categories or modalities of knowledge there are, and which has access to which other. It also depends upon the microstructure of the system -- how items are represented within each box, and how representations in one box activate representations in other boxes.

The question of whether PDP models accurately reflect the microstructure of human cognition is a controversial one, which cannot be settled on the basis of any single result. Nevertheless, the present results suggest that two very general properties of PDP models are explanatory of some otherwise puzzling phenomena, and hence provide some degree of confirmation for the psychological reality of at least these properties of PDP. The first property is the involvement of all parts of a network, directly or indirectly, in the computations that intervene between an input in one part of the system and an output in another part. This property accounts naturally for the effects of damage localized to one part of semantic memory on the ability to associate names and pictures of items that are represented in still-intact parts of semantic memory. At the macroscopic level of analysis, it is not clear why eliminating one of two or more possible routes from pictures to names (such as pictures to functional semantics to names) should result in impaired ability to associate pictures with their names, so long as another possible route (such as pictures to visual semantics to names) is still intact. The second of these properties is the need for collateral support, in activating one portion of a representation, from other parts of the same representation. This property accounts naturally for the effects of damage to visual semantics on the retrieval of functional information about living things. Again, at the macroscopic level of analysis, it is not clear why loss of knowledge of the appearance of something would affect the ability to access knowledge of its functions. Thus, the explanatory power of the model presented here depends on it having these properties of PDP models. The PDP mechanisms are not an incidental aspect of the model's implementation; they play a crucial explanatory role.

Table captions

Table 1. Performance of two patients with impaired knowledge of living things on various semantic memory tasks.

Table 2. Performance of two patients with impaired knowledge of nonliving things on various semantic memory tasks.

Table 3. Performance of the basic model, as measured by probability of correctly associating names and pictures, for living and nonliving things following different amounts of damage to visual semantics units and functional semantics units. Standard error of the mean shown in parentheses.

Table 4. Performance of the basic model, as measured by the dot product of the correct and obtained semantic patterns, for living and nonliving things following different amounts of damage to visual semantics units and functional semantics units. Standard error of the mean shown in parentheses.

Table 5. Performance of the basic model for functional knowledge of living and nonliving things, as measured by the dot product of the correct and obtained functional semantic patterns following different amounts of damage to visual semantics units. Standard error of the mean shown in parentheses.

Table 6. Performance of a patient whose semantic memory impairment was confined to knowledge of living things when probed verbally.

Table 7. Performance of the basic model, as measured by probability of correctly associating names and pictures, for living and nonliving things, probed verbally and pictorially, following different degrees of damage to the connections linking name units to visual units. Standard error of the mean shown in parentheses.

Table 8. Performance of the basic model, as measured by the dot product of the correct and obtained semantic patterns for living and nonliving things, probed verbally and pictorially, following different degrees of damage to the connections linking name units to visual units. Standard error of the mean shown in parentheses.



Figure captions

Figure 1. Schematic diagram of the model.

Figure 2. Performance of the basic model, as measured by probability of correctly associating names and pictures, for living and nonliving things following different amounts of damage to visual semantics units and functional semantics units.

Figure 3. Performance of the four variants of the basic model, as measured by probability of correctly associating names and pictures, for living and nonliving things following different amounts of damage to visual semantics units and functional semantics units. (a) Training stopped after 50 epochs; (b) Training continued for 200 epochs; (c) Trained without weight decay; (d) Equal numbers of visual and functional semantics units.

Figure 4. Performance of the basic model, as measured by the dot product of the correct and obtained semantic patterns, for living and nonliving things following different amounts of damage to visual semantics units and functional semantics units.

Figure 5. Performance of the four variants of the basic model, as measured by the dot product of the correct and obtained semantic patterns, for living and nonliving things following different amounts of damage to visual semantics units and functional semantics units. (a) Training stopped after 50 epochs; (b) Training continued for 200 epochs; (c) Trained without weight decay; (d) Equal numbers of visual and functional semantics units.

Figure 6. Performance of the basic model for functional knowledge of living and nonliving things, as measured by the dot product of the correct and obtained functional semantic patterns following different amounts of damage to visual semantics units.

Figure 7. Performance of the four variants of the basic model for functional knowledge of living and nonliving things, as measured by the dot product of the correct and obtained functional semantic patterns following different amounts of damage to visual semantics units. (a) Training stopped after 50 epochs; (b) Training continued for 200 epochs; (c) Trained without weight decay; (d) Equal numbers of visual and functional semantics units.

Figure 8. Performance of the basic model, as measured by probability of correctly associating names and pictures, for living and nonliving things, probed verbally and pictorially, following different degrees of damage to the connections linking name units to visual units.

Figure 9. Performance of the four variants of the basic model, as measured by probability of correctly associating names and pictures, for living and nonliving things, probed verbally and pictorially, following different degrees of damage to the connections linking name units to visual units. (a) Training stopped after 50 epochs; (b) Training continued for 200 epochs; (c) Trained without weight decay; (d) Equal numbers of visual and functional semantics units.

Figure 10. Performance of the basic model, as measured by the dot product of the correct and obtained semantic patterns for living and nonliving things, probed verbally and pictorially, following different degrees of damage to the connections linking name units to visual units.

Figure 11. Performance of the four variants of the basic model, as measured by the dot product of the correct and obtained semantic patterns for living and nonliving things, probed verbally and pictorially, following different degrees of damage to the connections linking name units to visual units. (a) Training stopped after 50 epochs; (b) Training continued for 200 epochs; (c) Trained without weight decay; (d) Equal numbers of visual and functional semantics units.

Acknowledgements

This research was supported by ONR grant NS0014-89-J3016, NIMH grant R01 MH48274, NIH career development award K04-NS01405, NIMH career development award MH00385, and NSF grant BNS 88-12048. We thank Robin Rochlin for assistance in data collection.

### References

- Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977). Distinctive features, categorical perception and probability learning: Some applications of a neural model. Psychological Review, 84, 413-451.
- Basso, A., Capitani, E., & Laiconi, M. (1988). Progressive language impairment without dementia: A case with isolated category specific semantic defect. Journal of Neurology, Neurosurgery and Psychiatry, 51, 1201-1207.
- Farah, M. J. (1991) Patterns of co-occurrence among the associative agnosias: Implications for visual object representation. Cognitive Neuropsychology, 8, 1-19.
- Farah, M. J., Hammond, K. H., Mehta, Z., & Ratcliff, G. (1989). Category-specificity and modality-specificity in semantic memory. Neuropsychologia, 27, 193-200.
- Farah, M. J. (1991), McMullen, P. A., Meyer, M. M. (1991) Can recognition of living things be selectively impaired? Neuropsychologia, 29, 185-193.
- Goodglass, H., & Baker, E. (1976). Semantic field, naming, and auditory comprehension in aphasia. Brain and Language, 3, 359-374.
- Goodglass, H., Wingfield, A., Hyde, M. R., & Theurkauf, J. C. (1986) Category-specific dissociations in naming and recognition by aphasic patients. Cortex, 22, 87-102.
- Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed Representations. In D. E. Rumelhart, & J. L. McClelland (Ed.), Parallel Distributed Processing: Explorations in the Microstructure of Cognition Cambridge: The MIT Press.

- McCarthy, R. A. & Warrington, E. K. (1988) Evidence for modality-specific meaning systems in the brain. Nature, 334, 428-430.
- McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. Journal of Experimental Psychology: General, 114, 159-188.
- Mehta, Z., & Newcombe, F. (1990). Paper presented at the 1990 Meeting of the International Neuropsychology Symposium. Nafplio, Greece.
- Pietrini, V., Nertimpi, T., Vaglia, A., Revello, M. G., Pinna, V., & Ferro Milone, F. (1988). Recovery from herpes simplex encephalitis: Selective impairment of specific semantic categories with neuroradiological correlations. Journal of Neurology, Neurosurgery, and Psychiatry, 51, 1284-1293.
- Sartori, G., & Job, R. (1988). The oyster with four legs: A neuropsychological study on the interaction of visual and semantic information. Cognitive Neuropsychology, 5, 105-132.
- Seidenberg, M. S. & McClelland, J. L. (1989) A distributed, developmental model of word recognition and naming. Psychological Review, 96, 523-568.
- Silveri, M. C., & Gainotti, G. (1988). Interaction between vision and language in category-specific semantic impairment. Cognitive Neuropsychology, 3, 677-709.
- Warrington, E. K., & McCarthy, R. (1983). Category specific access dysphasia. Brain, 106, 859-878.
- Warrington, E. K., & McCarthy, R. (1987). Categories of knowledge: Further fractionation and an attempted integration. Brain, 110, 1273-1296.

Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. Brain, 107, 629-854.

Table 1.

Case	Picture identification	
	Living	Nonliving
JBR	6%	90%
SBY	0%	75%

	Spoken word definition	
	Living	Nonliving
JBR	8%	79%
SBY	0%	52%

## Examples of definitions

## Living things:

- JBR Parrot - don't know  
 Daffodil - plant  
 Snail - an insect animal  
 Eel - not well  
 Ostrich - unusual
- SBY Duck - an animal  
 Wasp - bird that flies  
 Crocus - rubbish material  
 Holly - what you drink  
 Spider - a person looking for things, he was a spider for his  
 nation or country

## Nonliving things:

- JBR Tent - temporary outhouse, living home  
 Briefcase - small case used by students to carry papers  
 Compass - tools for telling direction you are going  
 Torch - hand-held light  
 Dustbin - bin for putting rubbish in
- SBY Wheelbarrow - object used by people to take material about  
 Towel - material used to dry people  
 Pram - used to carry people, with wheels and a thing to sit on  
 Submarine - ship that goes underneath the sea  
 Umbrella - object used to protect you from water that comes

Table 2.

Case	Spoken word-picture matching		
	Animals	Flowers	Objects
VER	86%	96%	63%
YOT	86%	86%	67%

Picture-picture matching		
	Animals	Objects
YOT	100%	69%



Table 3.

## Damage to visual semantic memory

Amount of damage (% of visual semantic units destroyed)	Probability correct (standard error of mean)	
	Nonliving things	Living things
0	1.00 (0)	1.00 (0)
20	0.97 (.02)	0.98 (.02)
40	0.91 (.04)	0.86 (.05)
60	0.88 (.05)	0.70 (.07)
80	0.80 (.06)	0.22 (.06)
99	0.73 (.06)	0.05 (.03)

---

## Damage to functional semantic memory

Amount of damage (% of functional semantic units destroyed)	Probability correct (standard error of mean)	
	Nonliving things	Living things
0	1.00 (0)	1.00 (0)
20	1.00 (0)	1.00 (0)
40	0.93 (.04)	1.00 (0)
60	0.88 (.05)	1.00 (0)
80	0.87 (.05)	1.00 (0)
99	0.73 (.06)	1.00 (0)

Table 4.

## Damage to visual semantic memory

Amount of damage (% of visual semantic units destroyed)	Scaled dot product in semantic units (standard error of mean)	
	Nonliving things	Living things
0	1.00 (0)	1.00 (0)
20	0.87 (.02)	0.84 (.02)
40	0.72 (.02)	0.69 (.02)
60	0.67 (.02)	0.59 (.02)
80	0.50 (.01)	0.40 (.02)
99	0.42 (.01)	0.32 (.01)

---

## Damage to functional semantic memory

Amount of damage (% of functional semantic units destroyed)	Scaled dot product in semantic units (standard error of mean)	
	Nonliving things	Living things
0	1.00 (0)	1.00 (0)
20	0.77 (.02)	0.77 (.02)
40	0.60 (.02)	0.65 (.02)
60	0.55 (.02)	0.61 (.02)
80	0.49 (.02)	0.51 (.02)
99	0.36 (.01)	0.40 (.01)

Table 5.

## Damage to visual semantic memory

Amount of damage  (% of visual semantic units destroyed)	Scaled dot product in functional semantic units (standard error of mean)	
	Nonliving things	Living things
0	1.00 (0)	1.00 (0)
20	0.95 (.02)	0.92 (.03)
40	0.91 (.02)	0.84 (.02)
60	0.89 (.02)	0.80 (.03)
80	0.84 (.02)	0.70 (.03)
99	0.81 (.02)	0.65 (.03)

Table 6.

## Identifications of pictures and words

	Living	Nonliving
Spoken words	33%	89%
Pictures	94%	98%

## Examples of identifications of living things

## Rhinoceros:

Spoken word -- animal, can't give you any functions

Picture -- enormous, weighs over one ton, lives in Africa

## Dolphin:

Spoken word -- a fish or a bird

Picture -- dolphin lives in water... they are trained to jump up and come out... In America during the war they started to get this particular animal to go through to look into ships

Table 7.

## Damage to connections from names to visual semantics

Amount of damage (% of name-visual semantic connections destroyed)	Probability correct (standard error of mean)			
	Nonliving things		Living things	
	Picture	Name	Picture	Name
0	1.00 (0)	1.00 (0)	1.00 (0)	1.00 (0)
20	1.00 (0)	0.98 (.02)	1.00 (0)	0.92 (.04)
40	1.00 (0)	0.92 (.04)	1.00 (0)	0.90 (.04)
60	1.00 (0)	0.76 (.06)	1.00 (0)	0.56 (.07)
80	1.00 (0)	0.66 (.07)	1.00 (0)	0.30 (.07)
99	1.00 (0)	0.80 (.06)	1.00 (0)	0.00 (0)

Table 8.

## Damage to connections from names to visual semantics

Amount of damage  (% of name-visual semantic connections destroyed)	Scaled dot product in semantic units (standard error of mean)			
	Nonliving things		Living things	
	Picture	Name	Picture	Name
0	1.00 (0)	1.00 (0)	1.00 (0)	1.00 (0)
20	0.98 (.01)	0.94 (.01)	0.98 (.01)	0.91 (.01)
40	0.96 (.01)	0.87 (.01)	0.96 (.01)	0.79 (.02)
60	0.94 (.01)	0.76 (.02)	0.94 (.01)	0.63 (.02)
80	0.92 (.01)	0.68 (.01)	0.92 (.01)	0.46 (.03)
99	0.90 (.01)	0.58 (.02)	0.90 (.01)	0.23 (.03)




Figure 2

# Basic model


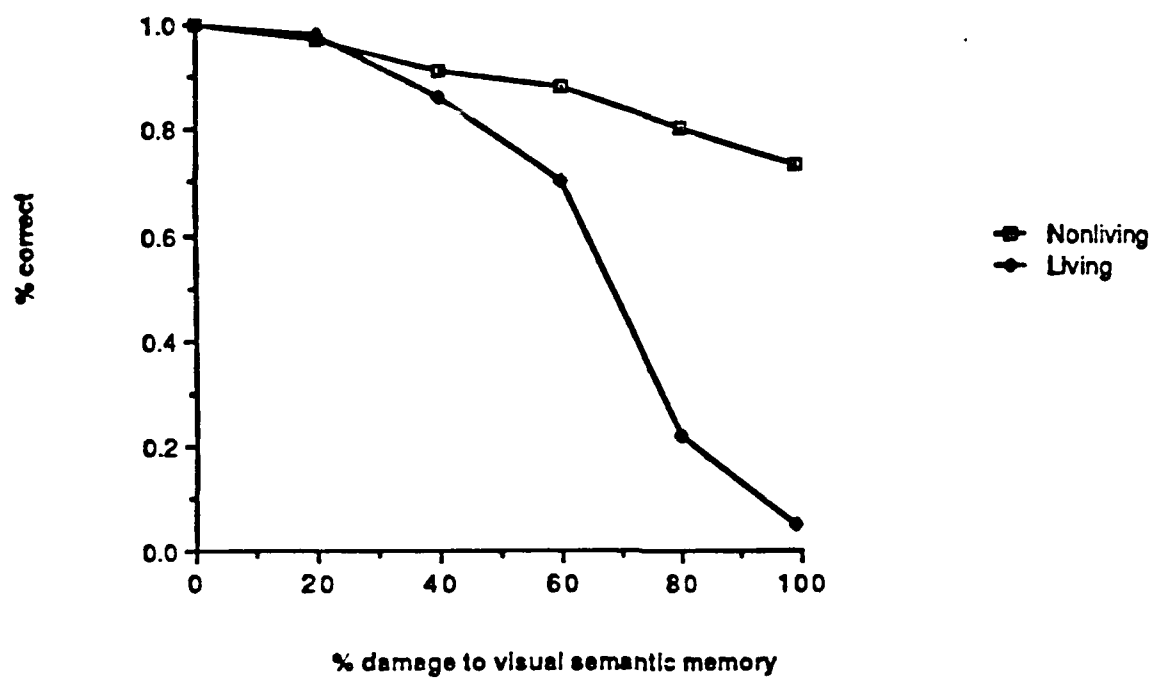




Figure 3a

Training stopped after 50 cycles




Figure 3b

Training continued for 200 cycles




Figure 3c

Trained without weight decay

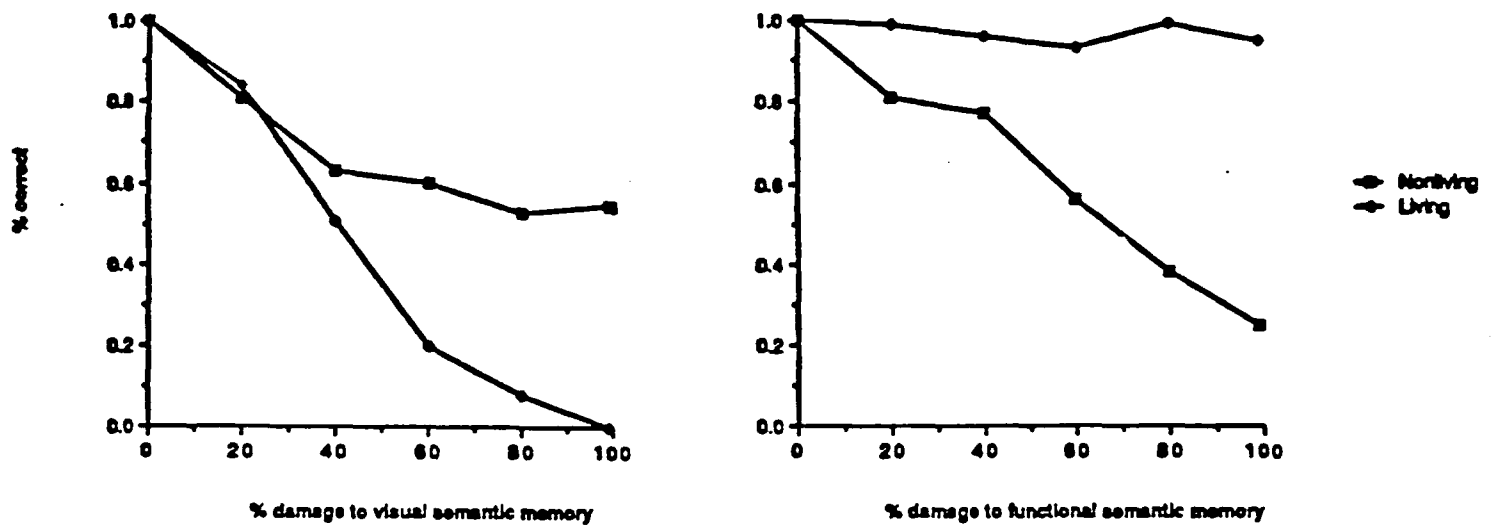


Figure 3d

Equal numbers of visual and functional semantic units




Figure 4

# Basic model





Figure 5a

Training stopped after 50 cycles




Figure 5b

Training continued for 200 cycles




Figure 5c

Trained without weight decay




Figure 5d

Equal numbers of visual and functional semantic units




Figure 6




Figure 7a  
Training stopped after 50 cycles




Figure 7b  
Training continued for 200 cycles




Figure 7c  
Trained without weight decay




Figure 7d  
Equal numbers of visual and functional semantic units




Figure 8






Figure 9a  
Training stopped after 50 cycles




Figure 9b  
Training continued for 200 cycles




Figure 9c  
Trained without weight decay




Figure 9d  
Equal numbers of visual and functional semantic units




Figure 10




Figure 11a

Training stopped after 50 cycles

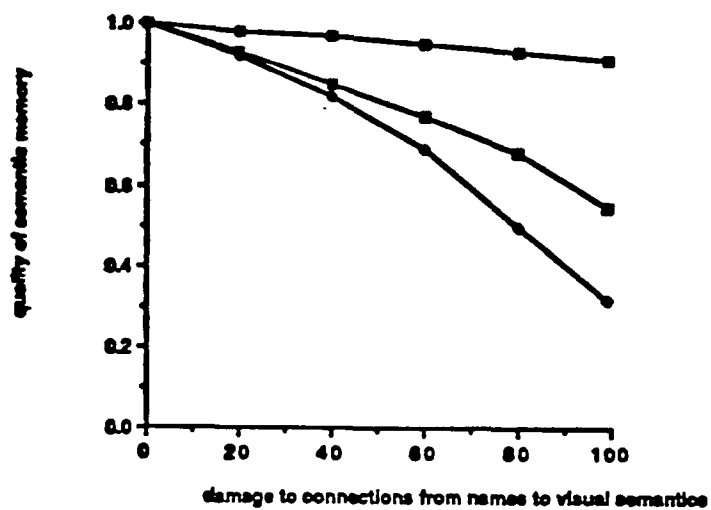


Figure 11b

Training continued for 200 cycles




Figure 11c

Trained without weight decay




Figure 11d

Equal numbers of visual and functional semantic units

